

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
3 October 2002 (03.10.2002)

PCT

(10) International Publication Number
WO 02/078253 A2

(51) International Patent Classification: **H04L 12/00**

(21) International Application Number: **PCT/GB02/01104**

(22) International Filing Date: **12 March 2002 (12.03.2002)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
0107638.9 27 March 2001 (27.03.2001) **GB**

(71) Applicant (for all designated States except US): **MARCONI COMMUNICATIONS LIMITED [GB/GB]**; New Century Park, P.O. Box 53, Coventry CV3 1HJ (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **HARDY, William, Geoffrey [GB/GB]**; 87 Monks Road, Binley Woods, Coventry CV3 2BQ (GB). **GRANDI, Vittoriano [GB/GB]**; 31 Station Avenue, Tile Hill, Coventry CV4 9HR (GB).

(74) Agent: **COLLIER, Ian, Terry**; Marconi Intellectual Property, Marrable House, The Vineyards, Great Baddow, Chelmsford, Essex CM2 7QS (GB).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.


(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: **ACCESS NETWORKS**



(57) Abstract: In order to enable access network to use private IP addresses, and so reduce public IP address overheads, data packets are tunnelled though the access network using one of a number of methods including, using layer 2 transmission protocol between a LAC and a LNS and using label switch paths based on MPLS labels.

WO 02/078253 A2

- 1 -

ACCESS NETWORKS

This invention relates to access networks for delivering IP services from telecommunications service providers to business and domestic customers.

5 Access networks are based on IP (Internet Protocol) and are a convenient way of delivering services to customers such as Video on Demand, telephony and multimedia. Such services may be delivered in a transparent manner. Each
10 Network Terminator (NT) in the access network may be provided with a number of service points such as, for example, management, voice over IP (VoIP), video services and Internet access. Each service point may be allocated an individual IP address. However, this construction is
15 wasteful of IPv4 addresses which are a relatively scarce resource and becoming more scarce.

Figure 1 shows an example of an access network 10 which connects a customer/user 12 to the Internet 14 via a router 16. The IP access network includes a number of network
20 terminators 18, each for delivering a specific service, and each having a unique public IP address. The network terminators 18 are all connected to a switch 11. In this example, the IP access network effectively forms part of the Internet. The IP addresses used within the access network
25 10 are all public IP addresses as are the addresses used by the end customers/users 12.

Router 16 between the access network 10 and the Internet 12 advertises its public network ID "Network A" to the rest of the Internet and all addresses within the access
30 network 10 are defined as hosts in network A.

The arrangement described is advantageous in principle but suffers from a number of disadvantages. First, the access network operator must obtain a large IP address space

- 2 -

from an Internet address allocation organisation. Some of these addresses will only be used for internal use within the access network while others will be used by users to connect to the Internet. This is a problem as IPv4
5 addresses are becoming scarce and it is undesirable to use more than the bare essential number of addresses.

The use of public addresses in the access network has potentially adverse security implications. These addresses, are, by definition, globally visible and the access network
10 operators may need to implement complex firewalls to provide adequate security. This is clearly expensive and so undesirable.

The number of IP addresses offered to each network terminator is fixed when the network is designed. The
15 network operator will usually want to minimise this number to conserve IP addresses. This makes it difficult for the network to provide for growth in the number of users. As a given customer adds more PCs to their network, there may come a time when the allocated IP addresses run out. This
20 problem can be dealt with by using Network Address Port Translation but it is not ideal as it runs contrary to the concept of ubiquity of public IP addresses in the whole network. Moreover, it can create problems with some IP protocols.

25 Despite the problems mentioned above, IP access networks are, in theory, desirable as they are simple and transparent to service provision. The invention aims to overcome the problems mentioned to make access network more practical to implement.

30 Accordingly there is provided a method of routing data packets from a client terminal to a destination through an access network, the access network having a network terminator, a plurality of network elements each having a private network address and a connection with a public
35 network, the method comprising tunnelling the data packets through the private access network to the connection with a public network.

- 3 -

The invention also provides a communications access network comprising a network terminator having a public and a private IP address and having a plurality of clients connected thereto, a plurality of network elements each having a private network address and a connection with a public network, and means for tunnelling data packets through the private access network from a client to the connection with a public network.

Embodiments of the invention have the advantage that by using tunnelling techniques to pass data packets across the access network, private addresses can be used for network elements in the access network. This enables the public IP address overhead to be reduced.

In one embodiment of the invention, the tunnel is an IP tunnel in which a private header is added to the packet to pass it through the access network.

In a further embodiment of the invention, the tunnel uses L2TP techniques with a LAC at either the client or the network terminator and an LNS at the other end of the tunnel. Data packets are passed between the LAC and LNS in PPP sessions.

In a third embodiment of the invention, the tunnel uses labels and sends the data packets along label switched paths. These labels may be MPLS labels.

Embodiments of the invention will now be described, by way of example, and with reference to the accompanying drawings, in which:

Figure 1 referred to above, is a schematic view of a prior art IP address network using public addresses throughout;

Figure 2 is a schematic view of an IP access network illustrating the principle of private addresses;

Figure 3 is a schematic view of an IP access network illustrating the principle of tunnelling and embodying the invention;

Figure 4 is a schematic view of a first embodiment of the invention;

- 4 -

Figure 5 shows how messages are sent from a client to a server via the IP access network in the embodiment of Figure 4;

5 Figure 6 shows how addresses are allocated in the embodiment of Figure 4;

Figure 7 is a schematic view of a first aspect of a second embodiment of the invention;

Figure 8 is a schematic view of a second aspect of the second embodiment of the invention;

10 Figure 9 is a schematic view of a third aspect of the second embodiment of the invention;

Figure 10 is a schematic view of a third embodiment of the invention and showing downstream labelling;

15 Figure 11 shows how labels are applied for upstream data flow;

Figure 12 shows an architecture to provide DHCP with MPLS in the embodiment of Figure 10;

Figure 13 shows automatic generation of labels in the embodiment of Figure 10;

20 Figure 14 shows how access MPLS tunnels and external MPLS tunnels can be integrated for upstream tunnels; and

Figure 16 shows how a single MPLS label can be allocated across a three stage network.

Referring to Figures 2 and 3, we have appreciated the
25 desirability of using private addresses within the access network. While this does not affect traffic within the IP access network itself, the IP access network is not transparent to external users. This may be overcome by using some sort of network address translation (NAT) at the
30 connection point between the IP access network and the Internet and at the connection point between the user and the IP access point.

Thus, in Figure 2, using the same numbering as in Figure 1, an address translator 20 is arranged between the
35 user 12 and the IP access network 10 and also between the IP access network 10 and the Internet 14. Thus, the Internet network address of the IP Access Network, Network A is

- 5 -

translated into a public network address p of the public network p.

Whilst this solution is adequate in theory, it remains problematic as many protocols do not pass transparently through network address translators. An application level gateway has to be added which processes all packets at the application layer to translate embedded IP addresses. Examples of such protocols include voice over IP (VoIP) protocols H.323 and SIP. The use of network address translation also prevents many common security protocols such as IPsec from being used. This is clearly unattractive to any security conscious user such as a business.

Despite the disadvantages mentioned, the solution outlined with respect to Figure 2 is attractive in that there is no overhead in the number of private addresses used in the access network. A large number of private addresses may be used and a small number of public addresses.

Figure 3 shows the principle behind the present invention. It retains the advantageous use of private addresses in the IP access network but does not use network address translation with its abundant disadvantages. Instead, the invention uses tunnelling techniques to offer end users public addresses. Thus, referring to Figure 3, tunnel end points 22, 24 are created between the user 12 and the IP access network 10 and the network 10 and the public network. All IP addresses used within the Access Network are within network A, but a user wishing to access the Internet is given an address from network P.

Thus, the user does not "see" the access network and datagrams are tunnelled through that network. The use of tunnelling is highly advantageous as the access network operator is free to choose the range of IP address without limitation as private addresses can be duplicated without adverse consequence.

There are no adverse security implications as the access network is not directly routable from the Internet;

- 6 -

the end user never sees the access network addresses. Furthermore, the tunnel acts as a transparent pipe, avoiding the problems highlighted with some protocols.

There are a number of tunnelling methods which are possible. The following will be considered: IP over IP; L2TP from NT with PPPoE or PPTP from user's PC to NT; L2TP from user's PC; and MPLS. It should be understood that other tunnelling techniques may be possible and are within the scope of the present invention the techniques mentioned will be considered in turn.

IP over IP

Referring to Figures 4 to 6, the tunnel stretches between the NT and an end point which includes a DHCP (Dynamic Host Control Protocol) server 26. When the host PC, user 12 boots up it requests an IP address by broadcasting a DHCP discover message to the network terminator 18. The network terminator encapsulates the message within the private IP domain and forwards it through the access network to the tunnel end point. This is done by adding a private header to the request which conceals the original source and destination from the private network.

The function of the DHCP server at the tunnel end point is to lease out public IP addresses. In practice, this may be a different server from that used by the access network to lease out private network addresses. The DHCP server intercepts the DHCP discover message and responds by offering a public address message, P.h. The response is tunnelled through the access network and arrives back at the host 14 which now knows all relevant information such as public IP address, default gateway, etc.

When the host 12 wants to send a datagram to a remote PC 28, it transmits the datagram to the network terminator 18 acting as the tunnel entrance. The network terminator encapsulates the datagram with a private IP address from network A and sends it through the access network to the tunnel endpoint. The original datagram is received and

- 7 -

transmitted into the Internet, where it is routed to destination PC 28.

5 A datagram from the Internet, for example from PC 28 intended for the user PC 12 also have to pass through the tunnel. This datagram will be received at the tunnel endpoint 24 between the public Internet and the private access network 10 with the destination address P.h.

10 The tunnel end point looks up the address to find the internal address allocated to the user PC 12 and encapsulates the datagram within an IP packet using private addresses from network A, and transmits it to the other end of the tunnel 24, at network terminals 18. The network terminator receives the packet, strips off the encapsulation and can then deliver the original datagram to the user PC
15 12.

As mentioned above, the tunnel entrance between the Internet 12 and the IP access network 10 includes a DHCP server 26. The tunnel end point keeps a table of external and internal addresses and performs the look-up operation to
20 find the relevant private address in the access network for a given public address. The DHCP server allocates the public IP addresses and the tunnel end point snoops on the actual public address allocated and adds it to its table against the private address. The operation of the DHCP server
25 will now be described with reference to Figure 5. In that figure, message flows are indicated by arrows with broad arrows, for examples arrow 30, indicating broadcast messages, and thin arrows, for example arrow 32, indicating unicast messages. Figure 5 shows four points in the
30 transmission path: the user 12 or client; the network terminator 18 which is also the tunnel entrance; the end of the tunnel 24; and the DHCP server.

In the following example, the first character (eg A or P) indicates the network number part of the IP address. The
35 following alphanumeric characters stand for the host part of the IP address. (Eg P.C stands for a PC and P.D stands for the DHCP server on network P). For convenience, the host

- 8 -

addresses of the various service points within the NT are indicated as N.x, where N is the NT number, and x is the number of the service points within it. Internal Access Network addresses are within network A, the private network, thus, the address of the tunnel ends in A.N.X. at the network terminator and A.E.0 at the tunnel end. Public addresses are within network P, so the user PC has the allocated address PC and the DHCP server address PD.

At step 100, the user 12 issues a DHCP discover message to the local broadcast IP address. This is a broadcast message including three parameters: the source, src=0, the destination, dest=broadcast and the MAC address of the client My MAC address. The MAC address is the user's own hardware address. The Network Terminator will receive the broadcast message and will recognise it as a DHCP request message. It will recognise that it has to tunnel it through the access network. At step 102, it encapsulates the message within an IP packet having a source src=A.N.3, the identity of the terminator 18, and at destination dest=A.E.0, the address of the end of the tunnel. The network terminator is configured with this address. This packet is sent through the private network as a unicast message.

At the end of the tunnel 24, the original DHCP discover message is received by stripping off the IP packet. The original message is then broadcast at step 104 on the local network. The DHCP server receives the message and, at step 106, allocates an external IP address to the user's hardware address and responds with a DHCP offer message. This reply is broadcast as the user does not yet know its IP address making a unicast inappropriate. The message sent has the following parameters: source src=P.D, the destination dest=broadcast, the MAC address of client 12, My MAC address, and the public address being offered to client 12 (IP=P.C). In the example described there is a single DHCP server. Multiple servers may be used in which case replies may be received from more than one server.

- 9 -

The tunnel endpoint 24 now receives the message from the DHCP server and at step 108 tunnels it to the network terminator in the same manner as before, adding an IP packet to the message. However, the tunnel entrance does not know to which network terminator the message should be sent. It should be recalled that messages are unicast though the IP access network and that there will be a number of network terminators (Figure 3). This problem is overcome either by keeping a record of outstanding DHCP discover messages at the tunnel endpoint 24 and where they have come from, and using this to form the destination address for the tunnel; or adding a tagged option to the discover message at the network terminator 18 or the tunnel endpoint 24. This tag is enclosed in the DHCP offer sent out to the user and contains the internal IP address of the network terminator which is used to direct the message through the IP access network to the correct network terminator. At this point it is stripped from the message together with the internal IP packet before the DHCP offer is sent to the user at step 110.

At step 112 the user 12 receives the DHCP offer and broadcasts a DHCP request. This is tunnelled to the DHCP server in steps 114 and 116 in the same manner as described. Where there is only a single DHCP server the message could be unicast. However, where there are multiple DHCP servers a broadcast message is necessary as it acts as a refusal to other DHCP servers that may have responded to the original DHCP discover message. The purpose of the DHCP request message is to indicate acknowledgment of the acceptance of the public IP address by the client. This request will identify the address of the DHCP server which sent the IP address that has been accepted.

Finally, at steps 118, 120 and 122 the DHCP server responds with a DHCP acknowledge message which is tunnelled through the IP access network to the user in the same manner as described and which contains additional configuration data. During the above DHCP sequence, the tunnel endpoint

- 10 -

sets up means, for example a translation table, to allow the translation of external IP addresses to internal IP addresses within the tunnel. This allows the messages from the DHCP server and data packets received from the external network to be tunnelled to the correct NT.

Referring now to Figure 6, the address allocation of an IP access network using IP tunnelling will now be described. The system shown includes three private address networks A, B and C. Private address network A is the IP access network 10 and networks B and C are private IP networks that may be used by the network provider to concentrate traffic from a number of IP access networks. Networks P, Q and R are public address networks. Network P is the network 12 referred to earlier and is used by the Internet Service Provider (ISP) to provide a service to clients of the access network. It is subtended to those clients at 34, on the left of the private address networks. Networks Q and R are part of the Internet.

Routers Rtr1 to Rtr5 are arranged between the various networks. Router Rtr1 advertises network A to network B, that is to Router Rtr2; Router Rtr2 advertises to network C, or router 3 that it has a route to network A. The advertising of private address stops here. Router 4 advertises network P to the rest of the Internet.

When a host computer 28 having the public address R.k on network R sends a datagram with destination P.h, that is the original user 12 of the earlier example, the datagram will be sent to its default router address R.I on router Rtr5. Rtr5 looks up network P in its routing table in standard manner and sends the datagram to the ISP's router Rtr4 on address Q.I.

Router Rtr4 will examine the datagram and discover that it has a source address equal to its own network address: P and will user ARP (Address Resolution Protocol) to find the MAC address corresponding to P.h, the public address of the destination PC.

- 11 -

At this point, router Rtr3, the tunnel end point router, must respond with its own MAC address. The datagram is then sent to router Rtr3.

Router 3 looks up the source address P.h in its
5 tunnelling table to find the address of the network terminator within the access network A.n. It encapsulates the original datagram within an IP packet with destination address A.n, looks up network A in its private network routing tables and forwards the message to Rtr2 on address
10 C1.

Router Rtr2 forwards the message to router Rtr1 which is at the head of the network. The datagram is then routed through the access network to the relevant network terminator having address A.n.

15 At the network terminator, the received message has the IP header stripped off to recover the original datagram. The datagram is then delivered in conventional fashion using ARP on the client network.

Upstream packets sent from the user P.h to PC R.k will
20 now be described.

User P.h is configured with address P.M as its default gateway. This is effectively the public address of the tunnel entrance. The user PC uses ARP to find the MAC address of the network terminator and then transmits the
25 datagram to it. All network terminators may have the same gateway address P.M.

The network terminator receives the datagram and encapsulates it within an IP datagram having destination address C.2, the private IP address of the end of the
30 tunnel. The network terminator needs prior knowledge of this address which could be configured during the setup of the access network, or chosen, for example from a web page offered by a http server in the network terminator. Different tunnel endpoint addresses may be chosen for
35 different IPS's although only one endpoint can be used at a time by all clients connected to an NT as it is not possible to signal session information to the NT.

- 12 -

The datagram is routed through the access network to the head end router Rtr1 through network B to router Rtr3, the tunnel endpoint router, on address C.2. Router Rtr3 removes the tunnel header and recovers the original datagram with destination address R5. It looks up network R in its public network routing table and routes the datagram to the required host via routers Rtr4 and Rtr5.

Tunnelling using layer 2 Tunnelling Protocol (L2TP)

The use of layer 2 tunnelling protocol to tunnel through the access network will be described with reference to Figures 7, 8 and 9. In many respects, the manner in which messages are handled is similar to the embodiments described previously and so will not be described in a great detail.

Layer 2 tunnelling protocol has been introduced to provide efficient dial-up access to the Internet. The present embodiment adapts that usage by removing the conventional dial up element to provide access to public IP addresses from a privately addressed Access Network.

In Figure 7, there are illustrated two methods in which L2TP is used to provide Internet access. Figure 7 shows an access network 10 to which hosts 12, 13 are connected through network terminators 18. The access network is connected to the Internet 14 through a router 16 and, through a series of further routers to a further host PC 28.

L2TP was conceived to tunnel PPP (Point to Point Protocol) sessions across an IP network. Tunnelling is between a L2TP Access Concentrator (LAC) at one end and an L2TP Network Server (LNS) at the other. Both the LAC and LNS are known components and their structure need not be discussed. As the protocol works by transporting clients' PPP sessions to the LNS it allows IP addresses to be allocated remotely at the LNS and transferred to the PC. It will be appreciated that this is similar to the allocation of IP addresses by the DHCP server in the previous embodiment.

- 13 -

The LAC may be located in the network terminator. In Figure 7 the terminator 18 connected to host H, 12 is shown with a LAC 37. The terminator will also include a PNS (Point to Point Network Server) or a PPoE server (Point to Point over Ethernet) 38 to handle communications with the host PC.

The PPP protocol provides the capability to transport IP addresses. Host, H, 12 initiates a PPP session with the PNS in the NT using Point to Point Tunnelling Protocol (PPTP) or with the PPoE server using Point to Point Protocol over Ethernet (PPPoE). The PNS or PPoE server in the NT causes the LAC within the NT to initiate a L2TP session with the LNS. When the L2TP tunnel has been created, the client's PPP session is extended to the LNS using the L2TP tunnel. The only internal IP address required is the internal address of the LAC. Multiple PCs connected to the Ethernet port of the network terminator can create separate sessions over the Ethernet and receive individual IP addresses from the LNS.

In addition, a DHCP server may be provided in the network terminator 18 to provide IP addresses local to the customer's LAN. The addresses are not used by the Access Network or the Internet.

The second variant is to use the clients PC as the LAC. PC 13 in Figure 7 is shown configured as the LAC. This is possible if the PC is running the Windows 2000 Operating System from Microsoft Corp. which provides support for L2TP. Any other operating system offering such support would be appreciated.

All client PCs connected to the network terminator's Ethernet port are allocated an IP address by the access network. This enables messages to be routed between the PC based LACs and the LNS. These IP addresses may be allocated from the Access Network private address space or a network address (NAT) function may be provided in the network terminator 18a and a separate address space provided for the client LAN using a DHCP server. This latter arrangement is

- 14 -

illustrated in Figure 8 with the NAT shown at 40 and the DHCP server at 42, both within NT 18a.

In Figure 8, there are three network addresses; network A, P and C. Network A is the private address space of the access network operator; network P is the public address of clients using the Internet; and network C is the private address within the client's own LAN.

The NAT 40 has an internal address A.n in the access network. The DHCP server 42 within the 18a allocates addresses for the client within the client network C. the NT itself has an address C.d in the client domain. Thus, the host G, 13 receives a network address C.g from the DHCP server 42. The NAT 40 translates addresses between client domain address C and access domain addresses A.

When a client uses the internal LAC to connect to the ISP, the LNS will allocate a public address from network P. this IP address is passed via L2TP to the client PC which appears to the Internet as a detached part of Network P.

Figure 9 shows a variant of the first of the L2TP methods described in that example, the PNS server and LAC are located at the network terminator 18. In Figure 9 these two components are arranged at a central point. As can be seen in Figure 9, this point is between the access network 10 and the Internet 14, specifically before the Internet router 16.

The PPP session is then tunnelled from the user's PC to the PNS server 38 using PPTP (Point to Point Tunnelling Protocol). The tunnel is from a point to point Concentrator (PAC) at the PC. In this case the PAC is used as the client end of the PPTP protocol. The PPP session is then extended to the user's ISP using L2TP. The user is then allocated a public IP address in the domain allocated to his chosen ISP, that is the network served by the LNS belonging to the ISP.

Tunnelling Using MPLS

Figures 10 to 16 show a third embodiment of the invention in which MPLS (multi-protocol label switching) is

- 15 -

used to tunnel data through the access network. Use of MPLS has a number of advantages, namely it can be used to determine the physical path through the network. Instead of using MAC or IP addresses to route packets, MPLS can be generated according to the destination of the packets. MPLS can also be used to identify the quality of service requirements of paths through the network and provide multiple paths through the access networks.

The use of MPLS will be described first by considering downstream and upstream tunnelling with reference, respectively, to Figures 10 and 11.

Figure 10 shows the access network 10 having a network terminator 18, and a pair of concentrators 11 and an access network router 15. An explicitly router ISP is used to tunnel downstream data through the network. The access router 15 keeps a map of IP addresses to MPLS labels. When a packet arrives at the access router, its IP address is examined. Three MPLS labels, D1, D2 and D3 are inserted into the packet and the packet sent to the first stage concentrator 11a. The number of labels attached will be equal to the number of stages in the network through which the packet has to pass. In this case, there are three stages; access router to concentrator 1; concentrator 1 to concentrator 2; and concentrator 2 to network terminator.

The first stage concentrator examines the label on top of the stack D1 and uses it to route the packet, removing that label, D1 from the label stack. D1 may contain the output part number on which the packet is to be transmitted. Label D1 is popped off the label stack and the packet forwarded to the second stage concentrator 11b. Here a similar operation is performed, using label D2 and, according to the destination given by label D2 the packet, now only containing the original packet and label D3 is forwarded to the network terminator. At the NT 18, a similar operation is performed again, with the NT examining the remaining label D3 and routing the bare packet to the appropriate element in the network terminator depending upon

- 16 -

the routing information contained in label D3. This final destination is the tunnel endpoint.

The MPLS labels can also be used to provide quality of service QoS management by using a part of the label to allocate a class to the traffic which controls the queuing algorithms used on concentration points.

The embodiment has been described in terms of a label for each stage of the routing through the IP access network. The MPLS label is a standard length of 20 bits and a single label can carry routing and QoS information for more than one stage. This will be described later.

Referring now to Figure 11, upstream routing of packets is more simple as they are all destined for the same point; the access router 15. Thus, a single label only is required and is used by all the stages. The label is not popped up by any of the stages but merely examined before the packet and label is passed on to the next stage. The label is only popped at the access network router. Again, the label, shown as .U (upstream) in Figure 11 can also include QoS management, using different label values for different traffic classes.

It will be appreciated from the discussion of Figure 10 and 11 that the access network does not use IP addresses for internal routing of user packets. IP addresses are only used at the extremities of the access network where it has to communicate with external networks, for example at the access router 15 and the network terminator 18. Individual address domains may be used for each type of service offered by the NT, such as videos, voice over IP and Internet access to simplify the provision of firewall security.

Figure 12 illustrates how DHCP can be provided with MPLS tunnelling. Like components are shown with the same reference numerals as in previous examples.

The host 12 will request an IP address by generating a DHCP discover message. This arrives at the MPLS tunnel entrance 22 in the network terminator 22. The request is sent along the upstream LSP to the access router 15 in the

- 17 -

manner described with respect to Figure 11 the access router here acts as the tunnel endpoint 24. The DHCP discover request will now be acted upon by the DHCP server 26 which will allocate a public IP address to the client and send
5 this back to the client. To enable this, the access server 15 sets up the necessary mapping from IP address to MPLS label and sends the DHCP offer message along the downstream LSP back to the client in the manner described with respect to Figure 10.

10 MPLS labels may be generated automatically. This will be described with reference to Figure 13. To begin with, a special MDLS label Ud is reserved for DHCP discover and request messages. The network terminator 18 detects the DCHP message as it is an IP Broadcast message.

15 Broadcast messages are not normally forwarded by the network terminator. The NT inserts the MPLS label Ud and inserts the port number on which the request was received into a reserved field in the DHCP message. In the Figure 13 example, this is 002 hex. The DHCP request is then
20 forwarded on to the second concentrator stage 11b.

As each concentration stage receives the message it will recognise that the message is a DCHP request as the packet will carry the unique Ud label. The concentration inserts the port number on which the request was received
25 into some bits of the reserved field and passes the message on. In the present example it can be seen that the message is received at port three of concentration 110 so the reserved field changes from 002 to 032. At the next concentrator the message is received at port 1 and so the
30 reserved field changes to 132.

When the DCHP message is received at the access router, acting as the tunnel endpart, the reserved field will contain the port numbers on which the message was received at all the concentrator stages including the network
35 terminator. The DHCP request is sent to the DHCP server 26 and, when a response is received, the reserved field, which must be echoed by the DHCP server, can be used to generate

- 18 -

MPLS routing labels for the downstream path from the access server 15 to the network terminator 18.

One field which may be used as the reserved field is the chaddr field. If unicast DHCP renewals are used by clients, the NT also has to detect such renewals as a special case in order that the correct MPLS label can be applied.

So far, MPLS tunnels have been described purely within access networks. Access tunnels may be integrated with external MPLS tunnels as will be described with reference to Figures 14 and 15. The purpose of such integration is to enable the QoS attributes of the external tunnel to be maintained in the access network.

Figure 14 illustrates how this may be achieved for downstream messages. Here there are two separate downstream tunnels, LSP1 and LSP2. In the first tunnel, a packet is sent from server 50 to the IP access network router 15. This packet has an attached label Li which includes quality of service management information. The access router 15 terminates the tunnel LSP1 and pops the label Li extracting the QoS management information and the destination and generates labels D1 to D3, or whatever labels are required as discussed with respect to Figure 10. The QoS characteristics of tunnel LSP1 can be carried into these new labels so that the appropriate queues are used to forward the packets within the access network.

In Figure 15, upstream tunnels are easily integrated by extracting the quality of service information specified in an upstream label U in the access network at the access network router 15 and inserting it into the label of the second tunnel LSP2 to maintain continuity. Thus the label in the IP zone has the same QoS data.

It was mentioned earlier that downstream messages, which include several labels need not necessarily use a separate label for each stage. Figure 16 shows how a single 20 bit label could be allocated in a three stage access network. In Figure 16, the two concentrator stages 11a, 11b

- 19 -

are identified as street node and distribution nodes respectively. The access router is connected to 16 street nodes, each of which are connected to 32 distribution nodes, giving a total of 512 distribution nodes. The distribution nodes are each connected to 48 NTs; a total of 24575 NTs. Each of the NTs is connected to 8 service points each of which can be provided with one of four levels of QoS. The 20 bit MPLS label is therefore made up of a 4 bit street number, a 5 bit street node port, a 6 bit distribution node port, a 3 bit NT port and a 2 bit QoS.

Trade offs may be made in the bit allocations. For example, 32 street nodes each parenting 16 distribution nodes could be supported by allocating 5 bit to the street node number and four bits to the street node port number. At present, a two bit QoS is sufficient as only four levels of QoS are used: video, voice, LAN data and management but the above allocation allows for eight for future use. The number of service points at the NT may be reduced to four, using 3 MPLS but, and the number of QoS levels reduced to 2, using a single MPLS bit. This releases two further bits to allow, for example, 32 street nodes to support up to 64 distribution nodes each.

It will be appreciated that in each of the embodiments described, tunnelling techniques have been used to send data through an access network which uses private internal address. Each of the tunnelling techniques allows data to pass through the private address network without the need to know those private addresses. This has the advantage of making it possible to construct access networks using private internal addresses so reducing the need to use scarce public IP addresses in such networks.

Variations and modifications to the embodiments are possible and will occur to those skilled in the art. For example, other tunnelling techniques may be possible beyond those exemplified. Such modifications are within the scope of the present invention.

- 20 -

Claims

1. A method of routing data packets from a client terminal to a destination through an access network, the access network having a network terminator, a plurality of
5 network elements each having a private network address and a connection with a public network, the method comprising tunnelling the data packets through the private access network to the connection with a public network.
2. A method according to claim 1, wherein the data
10 packets are tunnelled from the network terminator to the connection with a public network.
3. A method according to claim 2, comprising sending a message from the client to the public IP address of the network terminator, encapsulating the message in a IP
15 packet having the destination address of the connection with a public network and sending the message from the network terminator to the destination address.
4. A method according to claim 3, comprising a DHCP server at the connection with a public network, wherein
20 the message sent from the client is a A DHCP discover message, the method comprising removing the IP packet from the message received at the connection with a public network, sending the message to the DHCP server, sending a return message including a client identifier from the
25 server to the connection with a network, encapsulating the return message in an IP packet having the destination address of the network terminator, sending the return message to the network terminator, removing the IP packet from the return message at the network terminator and

- 21 -

sending the message from the network terminator to the client.

5. A method according to claim 4, wherein the connection with the external network maintains a record of outstanding DHCP discover messages received and their source address in order to route the reply message to the correct network terminator.

6. A method according to claim 4, wherein the DHCP discover message is tagged with the private address of the network terminator and the tag is stripped before the DHCP discover message is sent to the DHCP server.

7. A method according to claim 1 or 2, wherein the network terminator comprises a PNS or PPP over Ethernet server and a L2TP (Layer 2 Tunnelling Protocol) access concentrator LAC, and the connection with an external network comprises a L2TP Network Server LNS, the method comprising establishing a PPP session between the L2TP access concentrator and L2TP network server and using the PPP session to tunnel data through the private access network.

8. A method according to claim 7, wherein the L2TP network server allocates a public IP address to the client and forwards the address as a part of the PPP session to the L2TP access concentrator.

9. A method according to claim 1, wherein the network terminator comprises a DHCP server and an address translator, the client includes a L2TP access concentrator and the connection with an external network comprises a L2TP network server, the method comprising establishing a PPP session between the client L2TP access concentrator and the L2TP network server and using the PPP session to tunnel data through the private access network.

- 22 -

10. A method according to claim 1 or 2, wherein the client comprises a point to point access concentrator (PAC) and the connection with an external network comprises a point to point network server (PNS), wherein
5 data is tunnelled across the private access network in a PPP session between the PAC and the PNS.
11. A method according to claim 1 or 2, wherein tunnelling the data packets through the network comprises attaching at least one label to the data packets based on
10 the IP address of the connection with an external network, the label including routing information through the private access network, the data packets being routed via a label switched path based on the routing information the at least one label.
12. A method according to claim 11, wherein the at each
15 label is an MPLS label.
13. A method according to claim 11 or 12, wherein a label is attached for each point in the network through which the data packets pass.
14. A method according to claims 11, 12 or 13, wherein
20 the or each label includes quality of service information.
15. A method according to any of claims 11 to 14, wherein the connection with an external network includes a DHCP server and the method comprises sending a DHCP discover
25 message from the network terminator via a label switched path to the connection with an external network, forwarding the DHCP discover message to the DHCP server and allocating a public IP address to the client at the DHCP server.
16. A method according to claim 15, comprising mapping
30 the allocated published IP addresses of the client to at

- 23 -

least one label at the connection with an external network and sending a message from the DHCP server including the client IP address via a label switched path to the network terminator.

5 17. A method according to claim 16, wherein the network terminator removes the at least one label and forwards the message from the DHCP server to the client.

10 18. A method according to claim 16 or 17, comprising inserting the port number on which the DHCP message is received at each stage of the label switched path into a reserve field within the message, and generating routing labels for routing the message from the DHCP server to the network terminator from the port numbers in the reserved field.

15 19. A method according to any of the claims 11 to 18. Further comprising tunnelling data packets from a second network point on the public network using a label switched path, and, at the connection with the public network, removing a label attached to the data packets received
20 from the second destination point and extracting the ultimate IP destination address therefrom, and generating a fresh set of labels to enable the data to be sent to the network terminator via a further label switched path.

25 20. A communications access network comprising a network terminator having a public and a private IP address and having a plurality of clients connected thereto, a plurality of network elements each having a private network address and a connection with a public network, and means for tunnelling data packets through the private
30 access network from a client to the connection with a public network.

- 24 -

21. A communications access network according to claim 20, wherein the network terminator includes the means for tunnelling data.

5 22. A communications access network according to claim 20 or 21, wherein the tunnelling means comprises means for encapsulating a message from a client in an IP packet having the address of the connection with a public network and means for sending the encapsulated message to the connection with an external network.

10 23. Apparatus according to claim 20 or 21, wherein the network terminator comprises an Ethernet server and a LAC, and the connection with an external network comprises a LNS whereby data can be tunnelled across the private access network by establishing a PPP session between the
15 LAC and the LNS.

24. Apparatus according to claim 20, wherein the network terminator comprises a DHCP server and an address translator, the clients each include a LAC, and the connection with an external network comprises an LNS
20 whereby data packets are tunnelled through the private access network by establishing a PPP session between the LAC of a given client and the LNS.

25. Apparatus according to claim 20 and 21, wherein the clients each comprise a PAC and the connection with all
25 external network comprises a PNS, whereby data packets are tunnelled across the private access network in a PPP session between the PAC and the PNS.

26. Apparatus according to claim 20 or 21, comprising means at the network terminator for generating at least
30 one label from the IP address of the data packets, means for attaching the at least one label to the data packets, and means for routing the data packets and at least one

- 25 -

label to the other of the network terminator and the connection with an external network via a label switched path.

Fig.1.




Fig.2.





Fig.3.



2/6

Fig.4.





Fig.5.



3/6

Fig.6.




Fig.7.




Fig.8.




Fig.9.





Fig.16



5/6

Fig.10.




Fig.11.





Fig.12.



6/6

Fig.13.




Fig.14.




Fig.15.

